UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/821,431	04/09/2004	Robert E. Cypher	5181-96100	1240

35690 7590 08/02/2007

MEYERTONS, HOOD, KIVLIN, KOWERT & GOETZEL, P.C.
P.O. BOX 398
AUSTIN, TX 78767-0398

EXAMINER	
MEONSKE, TONIA L	

ART UNIT	PAPER NUMBER
2181	

MAIL DATE	DELIVERY MODE
08/02/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Application Number: 10/821,431
Filing Date: April 09, 2004
Appellant(s): CYPHER ET AL.

Stephan J Curran
P.O. Box 398
Austin, TX 78767-0398

For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 10 May 2007 appealing from the Office action mailed 9 November 2006.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

2005/0223203	Loh	11-2005
2001/0056531	McFarling	12-2001
6,427,206	Yeh et al	7-2002

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Claims 1-4, 6-8, 13-17, 19-21, and 26-27 are rejected under 35 U.S.C. 103(a) as being unpatentable over Loh (US Pat. Appl. Pub. 2005/0223203; herein referred to as "Loh".) in view of McFarling (US Pat. Appl. Pub. 2001/0056531; herein referred to as "McFarling".).

Regarding **independent claim 1**, Loh discloses *a branch prediction mechanism comprising: a first storage including a first plurality of locations for storing a first set of partial prediction information* [see Loh, Fig. 4, element 405 (e.g., "Branch Predictor 1")]*; a second storage including a second plurality of locations for storing a second set of partial prediction information* [see Loh, Fig. 4, element 405 (e.g., "Branch Predictor 2"); Examiner's note: Loh does not limit the invention to a particular type of predictor, however, McFarling states that global history could be used to index the predictors, thus it is inherent that global predictors would be utilized as the predictors. An example of a common global predictor is shown in McFarling, Fig. 4. N.B., a global predictor contains an array of saturating counters.]*; wherein said control unit is further configured to provide a prediction value based on corresponding partial prediction information in said selected locations of said first and said second storages* [see Loh, Fig. 4, element 410; Para 0020, lines 3-7; Examiner's note: Since Loh discloses the prediction of segments of a branch history, it is clear that the intermediate predictions are for the segments and thus are only partial predictions, thus the need for a mechanism to determine a final prediction which is disclosed by Loh in element 410 of Figure 4.].

Loh does not disclose *performing a first hash function on input branch information to generate a first index for accessing a selected location within said first storage and* performing *a second hash function on said input branch information to generate a second index for accessing a selected location within said second storage wherein said input branch information includes address information corresponding to a fetch address of a current branch instruction or control unit further configured to update said selected locations of said first and said second storages dependent on whether said prediction value yields an accurate branch prediction.*

McFarling does disclose *performing a first hash function on input branch information to generate a first index for accessing a selected location within said first storage and* performing *a second hash function on said input branch information to generate a second index for accessing a selected location within said second storage* [see McFarling, Para. 0026, lines 10-13 and Figure 14 and Para. 0087-0089; Examiner's note: McFarling discloses utilizing a hashing function to index global predictors. And the case of performing multiple hash functions in which multiple storages can be accessed is shown], *wherein said input branch information* [including] *address information corresponding to a fetch address of a current branch instruction* [see McFarling, Fig. 5, element 50; Para. 0026, lines 10-13 "...the branch instruction address...".]. McFarling also discloses a *control unit...further configured to update said selected locations of said first and said second storages dependent on whether said prediction value yields an accurate branch prediction* [see McFarling, Para. 0021, lines 6-9; Examiner's note: McFarling discloses tracking branch instructions which would involve updating the counters in a branch predictor. Furthermore, the primary goal of a branch predictor employing saturating counters, such as the well known counters proposed by Lee and Smith, would have been to update said counters based on branch execution.].

The advantage of using a hash function to index a global predictor would have been to improve the global prediction accuracy [see McFarling, Para. 0026, lines 13-15]. This advantage is desirable in the environment disclosed by Loh as improving the accuracy of a branch prediction would increase the performance of an entire processor by reducing the amount of mispredictions and in turn fetching errors. This advantage would have motivated one of ordinary skill in the art at the time of invention to utilize a hashing mechanism, as disclosed by

McFarling, to index a set of predictors as disclosed by Loh. Therefore, it would have been obvious to one of ordinary skill in the art at the time of invention to use a hashing mechanism to index a set of predictors with the goal of increasing prediction accuracy.

Regarding **claim 2**, Loh discloses the *prediction value* [providing] *a strongly/weakly taken/not taken branch prediction indication that is indicative of whether the current branch instruction is taken upon execution* [see Loh, Para. 0020, lines 3-7; Examiner's note: Loh discloses a final prediction being a function of the partial predictions, if said partial predictions are based on a commonly used n-bit saturating counter (commonly known at the time of invention), it is inherent that the final prediction based on said partial predictions would have been of the same format.].

Regarding **claim 3**, Loh discloses said input branch information [including] branch history information corresponding to an outcome of a number of preceding branch instructions [see Loh, Para. 0021].

Regarding **claim 4**, Loh discloses *said first hash function and said second hash function...configured to operate on a portion of said branch history information* [see Loh, Para. 0018, lines 9-11].

Claim 5 has been cancelled.

Regarding **claim 6**, Loh discloses said first hash function and said second hash function...configured to operate on a different subset of bits of said fetch address [see Loh, Para. 0018, lines 9-11 and figure 4: Loh teaches "The prediction history information may be accessed in segments by a number of intermediate branch prediction units 405" and it is shown in figure 4 that these segments are non overlapping and thus would be of a different subset of bits].

Regarding **claim 7**, Loh discloses *each of said first and said second sets of partial prediction information* [including] *a plurality of counter values each corresponding to a strongly/weakly taken/not taken branch prediction indication that is indicative of whether the current branch instruction is taken upon execution* [see Loh, Para. 0020, lines 3-7; Examiner's note: Loh discloses a final prediction being a function of the partial predictions, if said partial predictions are based on a commonly used n-bit saturating counter (commonly known at the time of invention), it is inherent that the final prediction based on said partial predictions would have been of the same format.].

Regarding **claim 8**, Loh and McFarling do not explicitly disclose the *control unit* [being] *further configured to use said prediction value to determine whether the current branch instruction is taken upon execution, wherein said prediction value is generated by summing respective counter values stored within said selected location within said first storage and said selected location within said second storage.*

However, McFarling discloses using multiple values of saturating counters (typically ranging between the value of 0-3 as common at the time) and performing an action on said counters to derive a final prediction. At the time of invention it would have been obvious that an addition or average (inherently containing addition) of the partial counter values would have resulted in the best prediction with the least amount of overhead.

Claim 12 has been cancelled.

Regarding **claim 13**, Loh discloses *a third storage including a third plurality of locations for storing a third set of partial prediction information and wherein said control unit is further configured to perform a third hash function on said input branch information to generate a third index for accessing a selected location within said third storage* [see Loh, Fig. 4, element 405 (e.g., "Branch Predictor *k*"); Examiner's note: Loh does not place a limit on how many intermediate predictors are possible in the invention.].

Claims 14-17, 19-21, and 26 are rejected as being the method performed by the apparatus in claims 1-4, 6-8, and 14, respectively.

Claims 18 and 25 have been cancelled.

Claim 27 is rejected as being the branch prediction mechanism performed by the apparatus in claim 1.

Claims 9-11 and 22-24 are rejected under 35 U.S.C. 103(a) as being unpatentable over Loh in view of McFarling, further in view of Yeh et al. (US Pat. No. 6,427,206; herein referred to as "Yeh".).

Regarding **claim 9**, Loh, McFarling and Yeh disclose the limitations as stated in **claim 9**.

Yeh further discloses *each of said first and said second sets of partial prediction information* [including] *a plurality of counter values each corresponding to a strongly/weakly agree/disagree indication that is indicative of whether said branch prediction hint bit embedded within said current branch instruction is to be used by said control unit* [see Yeh, Col. 6, lines 18-35; Examiner's note: Yeh discloses using a hardware branch predictor if confidence in the compiler hint is not strong, that is, if the compiler is below a threshold of assuredness a branch prediction is further utilized to determine whether the hint bit is correct (e.g., hint = taken, predictor = strongly taken) or if the bit is incorrect (e.g., hint = taken, predictor = strongly not-taken).].

The advantage of utilizing a hint bit in conjunction with a branch prediction scheme would have been to determine in advance how certain a processor will be on the taken/not-taken tendencies of a branch instruction prior to its execution based on profiling the instruction by a compiler (Col. 2, lines 5-7). This advantage would have been desirable in the invention of Loh and McFarling as it would have sped up execution for branches that have been profiled correctly as strongly taken or strongly not-taken. Furthermore, as admitted by the applicant, modern processors are known to support a hint bit encoded by the compiler, therefore, it would have

been obvious to one of ordinary skill in the art at the time of invention to account for the hint bit in modern processors in the invention of Loh and McFarling. This advantage would have motivated one of ordinary skill in the art at the time of invention to account for the prediction hint and allow a branch prediction unit to use the hint to its advantage in determining the correct path of a branch.

Regarding **claim 10**, Loh and McFarling disclose the limitations as stated in **independent claim 1**.

Loh and McFarling do not disclose *the control unit is further* [being] *configured to use said prediction value to control whether a branch prediction is performed in accordance with a branch prediction hint encoded within a current branch instruction.*

Yeh does disclose *the control unit is further* [being] *configured to use said prediction value to control whether a branch prediction is performed in accordance with a branch prediction hint encoded within a current branch instruction* [see Yeh, Col. 6, lines 18-35; Examiner's note: Yeh discloses the use of a prediction "hint" (Col. 2, lines 62-64) encoded into the instruction by the compiler. This hint is utilized in conjunction with a branch prediction scheme (Col. 6, lines 18-35) to determine the appropriate branching behavior.].

The advantage of utilizing a hint bit in conjunction with a branch prediction scheme would have been to determine in advance how certain a processor will be on the taken/not-taken tendencies of a branch instruction prior to its execution based on profiling the instruction by a compiler (Col. 2, lines 5-7). This advantage would have been desirable in the invention of Loh and McFarling as it would have sped up execution for branches that have been profiled correctly

as strongly taken or strongly not-taken. Furthermore, as admitted by the applicant, modern processors are known to support a hint bit encoded by the compiler, therefore, it would have been obvious to one of ordinary skill in the art at the time of invention to account for the hint bit in modern processors in the invention of Loh and McFarling. This advantage would have motivated one of ordinary skill in the art at the time of invention to account for the prediction hint and allow a branch prediction unit to use the hint to its advantage in determining the correct path of a branch.

Regarding **claim 11**, Loh, McFarling, and Yeh disclose the limitations as stated in **claim 10.**

Loh, McFarling, and Yeh do not explicitly disclose *the prediction value* [being] *generated by summing respective counter values stored within said selected location within said first storage and said selected location within said second storage.*

However, McFarling discloses using multiple values of saturating counters (typically ranging between the value of 0-3 as common at the time) and performing an action on said counters to derive a final prediction. At the time of invention it would have been obvious that an addition or average (inherently containing addition) of the partial counter values would have resulted in the best prediction with the least amount of overhead.

Claims 22-24 are rejected as being the method performed by the branch prediction mechanism in claims 9-11, respectively.

(10) Response to Argument

Arguments begin on page 6 and continue onto page 16 of the Appeal Brief mailed 10 May 2007.

Appellant on page 8, beginning in the first paragraph and ending in the last, argues, "Appellant asserts Loh merely teaches the branch predictor units 405 making predictions using information obtained from the segment registers 401, and the final predictor making a prediction based off the intermediate predictions. Appellant submits Loh does not teach disclose or suggest that the branch predictor units 405 have any storage capacity whatsoever. Moreover Appellant cannot find any reference in Loh, whether implied or otherwise, that each predictor 405 may be a storage including a plurality of locations for storing a set of predictions. Appellant again notes that Loh does not provide any specific detail with regard to the structure of elements 405, other than what is discussed above...Appellant submits **Loh does not teach or suggest** 'a first storage including a first plurality of locations for storing a first set of partial prediction information,' and 'a second storage including a second plurality of locations for storing a second set of partial predication information,' as recited in Appellant's claim 1."

It is noted that the Appellant on the bottom of the second paragraph on page 8 states, "the branch history information is stored in the prediction history registers 401 and not the branch predictor 405. In addition, as illustrated in FIG. 4, Loh discloses one segment register 401 per branch predictor unit 405." This appears to be admission by the Appellant that there is a first and second storage that stores prediction information. The further admission that there is a one to one correspondence between each individual segment register 401 and each individual branch

predictor unit 405 means that the branch predictor 405 does indeed have some kind of storage capacity. In light of the statements by the Appellant, Examiner is unable to understand how Loh does not teach a first and second storage that stores prediction information.

Appellant argues in the last paragraph of page 9, "Appellant asserts this is the GShare branch prediction mechanism described in Appellant's background section. Further, Appellant asserts the use of the multiple hash functions as shown in Fig. 14 of McFarling, is not a direct extension of the use of one hash function." The appellant later goes on to argue on the top of page 11, "Appellant notes the second hash function is operating on at least a portion of the result from the first hash function," finally concluding in the last full paragraph of page 11, "McFarling uses the hashing functions differently than Appellant...McFarling does, in fact, use the two hash functions to access one cache structure, by hashing different information for a different reason."

It is noted that McFarling is used to teach the obviousness of using hashing functions and that it is unreasonable to believe that McFarling and Loh can be combined without adaptations. Loh teaches multiple cache structures that can utilize hashing functions and thus to properly combine Loh and McFarling, one having ordinary skill in the art would utilize more than one hash function. The use of two XOR's (as shown in figure 14 and paragraphs 0087-0089 of McFarling) is a simplistic example of a first and second hash function. It is noted that nowhere in the claims does it mention the complexity of the hash function, or whether the two hash functions are different. The contention of the Examiner, simply stated is McFarling teaches hash functions. McFarling also teaches using multiple hash functions, and claim 1 only requires "a first hash function on input branch information to generate a first index for accessing a selected

location within said first storage and to perform a second hash function on said input branch information to generate a second index for accessing a selected location within said second storage," which is taught by McFarling. The difference in the use of the hash functions, as argued by the Appellant, does not appear to be a claim limitation and thus Loh in view of McFarling teaches the claim limitations being argued.

Appellant argues in the last paragraph of page 12, "Appellant asserts Loh ONLY teaches 'a final branch history predictor unit 410 to generate a final branch prediction as function of the intermediate branch predictions performed by the intermediate branch prediction units.' Loh is silent as to how this is performed."

Although the rejection of claim 2 only recites paragraph 20 of Loh, it is noted that paragraphs 21 and 22 teach various other embodiments in which predictions are done in different manners. Loh teaches "the branch history register may contain a combination of various branch history information" (See paragraph 21). It is noted from the final office rejection that the Examiner has contended that the use of strength in branch prediction is well known. Various branch history information is able to read onto strength of branch predictions because such determinations must utilize branch history information.

Appellant argues in the first paragraph of page 13, "With regard to claim 8, the Examiner asserts that Loh discloses "using multiple values of saturating counters (typically ranging between the value of 0-3 as common at the time) and performing an action on said counters to derive a final prediction. Appellant can find absolutely no teaching of this feature in **Loh**."

It is noted that the rejection of claim 8 utilizes teachings from McFarling (See page 9 of the final rejection mailed 9 November 2006), which is not argued by Appellant. Appellant contends that Loh does not teach the claim limitations of claim 8, which the Examiner agrees with, however, the Examiner contends McFarling does as stated in the final rejection.

Appellant argues in the first full paragraph of page 15, "Appellant submits Yeh is teaching using the HW branch predictor if the hint confidence is not strong enough. This is not the same as "a plurality of counter values each corresponding to a strongly/weakly agree/disagree indication that is indicative of whether said branch prediction hint bit embedded within said current branch instruction is to be used by said control unit." Appellant goes on to argue in the fourth paragraph, "Appellant respectfully disagrees that it would be obvious to generate "said prediction value...by summing respective counter values stored within said selected location within said first storage and said selected location within said second storage" as recited in claims 8 and 11."

Appellant appears to be stating disagreement without much substance directed to an argument. It is also noted that Yeh is used to teach the concept that confidence in branch prediction was used by one having ordinary skill in the art at the time the invention was made.

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

07/24/2007

Tonia L. Meonske

Conferees:

Alford Kindred

Lynne Browne
APPEAL PRACTICE SPECIALIST, TQAS
TECHNOLOGY CENTER 2100